arXiv:1606.03472v1 [cs.IT] 13 May 2016

SUPER-RESOLUTION FROM BINARY MEASUREMENTS WITH UNKNOWN THRESHOLD

Subhadip Mukherjee, Anjany Kumar Sekuboyina, and Chandra Sekhar Seelamantula

Department of Electrical Engineering, Indian Institute of Science, Bangalore 560012, India
Emails: subhadip@ee.iisc.ernet.in, anjanykumar@outlook.com, chandra.sekhar@ieee.org

ABSTRACT

We address the problem of super-resolution of point sources from binary measurements, where random projections of the blurred measurement of the actual signal are encoded using only the sign information. The threshold used for binary quantization is *not known* to the decoder. We develop an algorithm that solves convex programs iteratively and achieves signal recovery. The proposed algorithm, which we refer to as the *binary super-resolution* (BSR) algorithm, recovers point sources with reasonable accuracy, albeit up to a scale factor. We show through simulations that the BSR algorithm is successful in recovering the locations and the amplitudes of the point sources, even in the presence of significant amount of blurring. We also propose a framework for handling noisy measurements and demonstrate that BSR gives a reliable reconstruction (correspondingly, reconstruction signal-to-noise ratio (SNR) of about 22 dB) for a measurement SNR of 15 dB.

***Index Terms*—** Super-resolution, binary compressive sensing, iteratively reweighted ℓ_1 minimization.

1. INTRODUCTION

Studying the finer details of a signal, beyond the resolution provided by measurement systems, is of importance to many applications such as medical imaging [1], astronomy [2], super-resolution localization microscopy [3–6], etc. The vast body of literature on super-resolution deals with the problem of recovering fine details of a signal from coarser measurements. In most super-resolution techniques, one essentially extrapolates the spectra of the signals of interest to higher frequencies, given their low-frequency measurements. As a concrete example, consider astronomical imaging, where one is interested in identifying and super-resolving celestial objects that appear as point sources in the images (often blurred) captured by telescopes. For several applications, it would be desirable to identify close-by sources, surpassing the barrier posed by the diffraction limit of the imaging device.

Super-resolution of point sources has been recently studied by Candès and Granda, for noiseless [7] and noisy [8] measurements. They addressed the problem of extrapolating the spectrum of a superposition of point sources up to a frequency f_{hi}, from its low-pass measurements up to a frequency f_{lo}, and showed that a convex optimization-based approach produces a stable estimate, such that the reconstruction error is proportional to the noise level times the square of $\frac{f_{\text{hi}}}{f_{\text{lo}}}$, provided the sources are separated by at least $\frac{2}{f_{\text{lo}}}$.

An important aspect that one needs to consider while building an acquisition system is the precision of measurement, dictating the cost, speed of acquisition, and the subsequent storage requirement. From a practical viewpoint, it is important to design a signal recovery algorithm that takes into account the effect of measurement quantization. It is a worthwhile exercise to extend the limits of the reconstruction techniques where the signal is measured with finite precision. Designing compressive imaging systems with one-bit measurements was considered by Bourquard et al. [9, 10], who proposed an acquisition hardware and reconstruction algorithm for image recovery starting from binary measurements. Our work focuses on developing an algorithm for the separation of point sources in the extreme case where the imaging device employs one-bit quantization in comparison against an *unknown threshold* and records only binary data, potentially leading to fast and inexpensive hardware. Our reconstruction approach is iterative and based on convex programming, yielding accurate estimates up to a scale factor.

Sparse recovery using binary measurements was first addressed by Boufounous and Baraniuk [11], who proposed a fixed-point continuation algorithm, treating the measurements as sign constraints and restricting the search space to the unit sphere to avoid degenerate solution. Gupta et al. [12] proposed an adaptive algorithm for support recovery from sign measurements, requiring a sample complexity of $\mathcal{O}\left(s \log l_x\right)$, where s and l_x denote the sparsity and signal dimension, respectively. Plan and Vershynin developed a linear programming-based sparse recovery algorithm from binary measurements [13], extendable to approximately sparse vectors. They also established a connection between one-bit compressive sensing (CS) and sparse logistic regression and developed a convex program-based approach robust to sign-flips [14]. Aside from the convex-relaxation-based approaches, greedy algorithms such as matching sign pursuit [15], adaptive outlier pursuit [16] and iterative algorithms such as binary iterative hard thresholding [17], randomized fixed-point iterations [18], etc. are also available in the literature. The suitability of the one-bit CS paradigm over its multibit counterpart, for several practical applications, and its robustness to non-linear distortions were studied in [19] and [20], respectively.

The threshold used for the one-bit quantizer in the context of CS plays a crucial role in the reconstruction performance. Adaptive algorithms for selecting the threshold have been developed by Kamilov et al. [21] and Fang et al. [22]. In the following, we highlight our contributions and indicate the major differences between our approach and the related algorithms.

1.1. Our Contribution

We address the problem of recovering a combination of point sources using a setup that acquires binary measurements, where only the sign information of the random linear projections of a blurred measurement of the actual source is recorded. The threshold

used for collecting the binary measurements is not known to the decoder/reconstruction algorithm unlike the frameworks considered in [21, 22]. In fact, our reconstruction algorithm also gives an estimate of the threshold jointly with the underlying signal starting from the binary measurements. We formulate the problem of signal recovery, albeit up to a scale factor, as a convex optimization, and develop a *binary super-resolution* (BSR) algorithm, leveraging the idea of *iteratively reweighted* ℓ_1 *minimization* (IRℓ_1) [23]. We demonstrate successful application of the BSR algorithm in recovering point sources and demonstrate its robustness to noise.

2. PROBLEM STATEMENT AND THE BSR ALGORITHM

Consider a sparse signal $x(n) = \sum_{j=1}^{s} \alpha_j \delta(n - n_j)$, consisting of s Kronecker impulses of amplitudes α_j, located at time instants n_j. Suppose $x(n)$ is acquired through a measurement device having a low-pass impulse response $h(n)$, giving rise to the measured signal $z(n) = \sum_{j=1}^{s} \alpha_j h(n - n_j)$. Instead of directly recording $z(n)$, we consider a measurement setup where one computes linear projections of $z(n)$ using random sensing signals $p_i(n)$, $1 \leqslant i \leqslant m$, compares them with a threshold τ, and encodes them as $+1$ or -1, depending on whether they exceed τ or not, respectively. More precisely, the measurement device collects binary measurements of the form $y_i = \text{sgn}(\langle z(n), p_i(n) \rangle - \tau), 1 \leqslant i \leqslant m$, where $\langle \cdot, \cdot \rangle$ indicates the inner-product, and sgn is the signum function. The encoder chooses a value of τ such that not all measurements are mapped to $+1$ or -1. In fact, the encoder could make an optimal choice of τ according to some predefined criterion. The measurement process is expressed concisely as

$$y_i = \text{sgn}(A\mathbf{z} - \tau)_i = \text{sgn}(AH\mathbf{x} - \tau)_i, \tag{1}$$

where $\mathbf{x}$ denotes the vector representation of $x(n)$, H denotes the linear convolution matrix corresponding to the impulse response $h(n)$, and the i^{th} row of A contains the random sensing signal $p_i(n)$. Given binary measurements as in (1), our objective is to recover the underlying signal $\mathbf{x}$, or, equivalently, to estimate the amplitudes and the locations of the Kronecker impulses.

2.1. Development of the BSR Algorithm

Defining $\Phi = AH$, the binary measurements in (1) are written as $y_i = \text{sgn}\left(\phi_i^T \mathbf{x} - \tau\right)$, where ϕ_i denotes the i^{th} row of Φ. We seek to recover a sparse vector $\mathbf{x}$, with the least number of non-zero entries, together with an appropriate τ, such that the recovered $(\hat{\mathbf{x}}, \hat{\tau})$ pair is consistent with the measurements. This is accomplished by solving

$$(\hat{\mathbf{x}}, \hat{\tau}) = \arg\min_{\mathbf{x},\tau} \|\mathbf{x}\|_0 \text{ subject to } y_i \left(\phi_i^T \mathbf{x} - \tau\right) \geqslant 0, 1 \leqslant i \leqslant m. \tag{2}$$

The inequality constraints in (2) ensure that the recovery is consistent with the sign measurements. Although the optimization posed in (2) captures the objective of sparse super-resolution from sign measurements, there are two major defects associated with it, namely, non-convexity and degeneracy. Non-convexity arises because of the ℓ_0-norm objective, whereas degeneracy is a consequence of the fact that $\hat{\mathbf{x}} = 0, \hat{\tau} = 0$ are feasible and yield the minimum possible value of the objective. However, it is possible to eliminate either defects as explained in the following.

Algorithm 1 BSR algorithm to estimate a superposition of point sources $\mathbf{x}$ from binary measurements $\mathbf{y}$ as in (1).

1. Input: The matrix $\Phi = AH$ and the iteration count I.

2. Initialization: Set $p \leftarrow 0$, $\lambda_i^{(p)} = 1$, $1 \leqslant i \leqslant l_x$, where l_x denotes the length of $\mathbf{x}$.

3. Iterate the following steps I times:

1. Obtain $\left(\hat{\mathbf{x}}^{(p)}, \hat{\tau}^{(p)}\right)$ by solving (5).
2. Update $\lambda_i^{(p)} = \frac{1}{\left|\hat{\mathbf{x}}_i^{(p-1)}\right| + \epsilon}$, for all i, and set $p \leftarrow p + 1$.

4. Scaling: Compute $\hat{\mathbf{x}} = \frac{\hat{\mathbf{x}}^{(I)}}{\left\|\hat{\mathbf{x}}^{(I)}\right\|_2}$.

5. Output: Estimated signal $\hat{\mathbf{x}}$.

2.1.1. Eliminating degeneracy and non-convexity

Since $y_i = \text{sgn}\left(\phi_i^T \mathbf{x} - \tau\right)$, the measurement consistency requirements can be expressed as

$$y_i \left(\phi_i^T \mathbf{x} - \tau\right) > \delta, 1 \leqslant i \leqslant m, \tag{3}$$

where $\delta = \min_{1 \leqslant i \leqslant m} y_i \left(\phi_i^T \mathbf{x} - \tau\right) \geqslant 0$. If the underlying signal $\mathbf{x}$ follows a continuous distribution, we have that $\delta > 0$, with probability one. Dividing both sides of (3) by δ and absorbing δ in $\mathbf{x}$ and τ, the inequalities in (3) can be rewritten in an equivalent form, given by $y_i \left(\phi_i^T \mathbf{x} - \tau\right) \geqslant 1$, for all i. Thus, one can avoid degenerate solution by rewriting (2) as

$$(\hat{\mathbf{x}}, \hat{\tau}) = \arg\min_{\mathbf{x},\tau} \|\mathbf{x}\|_0 \text{ subject to } y_i \left(\phi_i^T \mathbf{x} - \tau\right) \geqslant 1. \tag{4}$$

Although the solution to (4) is not degenerate, it suffers from scale ambiguity. To illustrate further, consider a solution $(\hat{\mathbf{x}}, \hat{\tau})$ to (4). For any $\beta > 1$, $(\beta\hat{\mathbf{x}}, \beta\hat{\tau})$ is also a solution, since it is feasible and yields the same value of the objective. To avoid this ambiguity, we assume that $\|\mathbf{x}\|_2 = 1$, and this condition is enforced after computing a solution to (4).

In order to avoid non-convex optimization, we employ IRℓ_1 to approximate the solution to (4). In the p^{th} iteration, we solve

$$\left(\hat{\mathbf{x}}^{(p)}, \hat{\tau}^{(p)}\right) = \arg\min_{\mathbf{x},\tau} \left\|\Lambda^{(p)}\mathbf{x}\right\|_1 \text{ s.t. } y_i \left(\phi_i^T \mathbf{x} - \tau\right) \geqslant 1, \tag{5}$$

where $\Lambda^{(p)}$ is a diagonal weight matrix, initialized with $\lambda_i^{(0)} = 1, \forall i$, and modified as $\lambda_i^{(p)} = \frac{1}{\left|\hat{\mathbf{x}}_i^{(p-1)}\right| + \epsilon}$, for every p. A constant $\epsilon > 0$ is added to the denominator for numerical stability. It is shown in [23] that the IRℓ_1 algorithm arises out of minimizing an objective of the form $f_0(\mathbf{x}) = \sum_{i=1}^{l_x} \log(|x_i| + \epsilon)$. The iterations are repeated until $p = I$, for some predetermined I. A convex optimization software, namely CVX [24], is employed to solve (5). The BSR procedure is summarized in Algorithm 1.

2.2. Illustration Using One Dimensional (1-D) Signals

To demonstrate the performance of the BSR algorithm, we consider a signal $x(n)$ of length $l_x = 200$, containing six impulses, as shown in Fig. 1(a). The signal is measured using a device having a low-pass response $h(n)$, which is assumed to be a sinc function, shown in Fig. 1(b). The low-pass measurement $z(n)$ of length $l_z = 300$ and

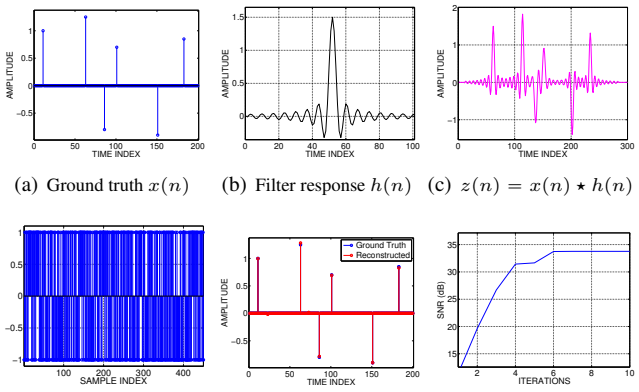

(a) Ground truth $x(n)$ (b) Filter response $h(n)$ (c) $z(n) = x(n) \star h(n)$

(d) Binary Measurements (e) Reconstructed signal (f) Recovery SNR

Fig. 1. (Color Online) Performance of BSR on 1-D signals.

the binary measurements y_i, $1 \leqslant i \leqslant m$, are plotted in Figures 1(c) and 1(d), respectively. We acquire $m = 1.5l_z = 450$ binary measurements of the form given in (1), where $\tau = -0.1$ and the sensing signals $p_i(n)$ are drawn independently, such that each entry of them takes a value of $+1$ or -1 with equal probability. The reconstructed signal, obtained using the BSR algorithm, and the variation of the signal-to-noise ratio (SNR), as a function of iterations, are shown in Figures 1(e) and 1(f), respectively. Using the measurement model (1), one requires to spend only $b = 1.5l_z$ bits for encoding the measurements, whereas, directly encoding $z(n)$ with a precision of 16 bits for every entry requires $b = 16l_z$ bits of storage. Therefore, the BSR algorithm results in a significant reduction in storage, without compromising on the reconstruction accuracy. The idea of using IRℓ_1 improves the SNR as the iterations progress, and the reconstructed signal after 10 iterations matches closely with the ground truth, as can be seen in Fig. 1(e).

3. EXPERIMENTAL RESULTS

We carry out experiments on images containing point sources to assess the efficacy of the BSR algorithm. A patch-based implementation of the BSR algorithm is developed in order to achieve scalability for images of larger sizes. We also conduct an empirical study to evaluate the performance of BSR under noise.

3.1. Patch-Based Implementation of BSR for Images

Let $\mathbf{x} \in \mathbb{R}^{l_x \times l_x}$ and $\mathbf{z} \in \mathbb{R}^{l_z \times l_z}$ denote the actual and the blurred images, where $\mathbf{z}$ is obtained by applying a low-pass mask h of size $p \times p$ on $\mathbf{x}$. For convenience, we assume that p is odd. Consider a patch of size $d \times d$, located at (i, j) in the blurred image $\mathbf{z}$. Using the vectorized representation $z^{(i,j)} \in \mathbb{R}^{d^2 \times 1}$ of the patch, we write

$$z^{(i,j)} = \mathcal{H} x^{(i,j)}, \tag{6}$$

where $x^{(i,j)} \in \mathbb{R}^{(d+p-1)^2 \times 1}$ denotes the vectorized version of the patch $x\left(i - \frac{p-1}{2} : i + d - 1 + \frac{p-1}{2}, j - \frac{p-1}{2} : j + d - 1 + \frac{p-1}{2}\right)$ of size $(d + p - 1) \times (d + p - 1)$ in the actual image x, and $\mathcal{H} \in \mathbb{R}^{d^2 \times (d+p-1)^2}$ is a matrix that captures the effect of the mask h on x. Binary measurements of the form

$$y_k^{(i,j)} = \text{sgn}\left(a_k^T \mathcal{H} x^{(i,j)} - \tau\right), 1 \leqslant k \leqslant m, \tag{7}$$

are acquired for each patch (i, j) in $\mathbf{z}$ without overlap, where a_k denotes a random vector of size d^2, for every k. The number of measurements for every patch is chosen as $m = 2d^2$. The threshold τ is unknown to the BSR algorithm and is estimated from the measurements. An estimate of the ground truth $\mathbf{x}$ is formed by stitching the recovered patches $\hat{x}^{(i,j)}$, considering the overlap of $(p - 1)$ pixels. The number of iterations in the BSR algorithm is taken as $I = 5$.

3.2. Performance Evaluation on Point-Source Images

Images of size 256×256, consisting of a superposition of $s = 100$ point sources are considered for experimental validation. The locations of the sources are selected uniformly at random from 256^2 possible locations and the magnitudes are drawn from a uniform distribution over $[0, 5]$. Gaussian blur kernels of different sizes p and standard deviations σ_h are applied on the ground truth image to obtain a low-pass measurement. Finally, random projections are computed for non-overlapping patches of size 16×16 extracted from the blurred image, followed by one-bit quantization. The number of binary measurements recorded for each patch is $m = 2 \times 16^2 = 512$. To reconstruct the ground truth image, the BSR algorithm is iterated $I = 5$ times. The performance is measured in terms of the ability to identify the locations and the amplitudes of the point sources.

Although the knowledge of the number of impulses s is not assumed in the BSR algorithm, we retain s dominant impulses in the recovered image for the purpose of comparison with the ground truth. To measure the accuracy of BSR in recovering the locations of the impulses, we evaluate the true positive ratio (TPR), defined as TPR $= \frac{|S \cap \hat{S}|}{|S|}$, where S and $\hat{S}$ denote the sets of locations of the impulses in the ground truth and the estimated images, respectively. The set $\hat{S}$ is computed after picking the s most significant impulses in the recovered image, ignoring the ones with negligibly small amplitudes. The TPR metric indicates the fraction of impulses whose locations are identified correctly. The amplitude recovery performance is measured by calculating the SNR over the correctly estimated impulses, denoted as SNR_1, and defined by $\text{SNR}_1 = \frac{\|\mathbf{x}_{S \cap \hat{S}}\|_F^2}{\|\mathbf{x}_{S \cap \hat{S}} - \hat{\mathbf{x}}_{S \cap \hat{S}}\|_F^2}$, capturing the discrepancy in amplitude estimation over the rightly identified impulses. We compute another metric referred to as the *relative error* (RE), defined as RE $= \frac{\|\hat{\mathbf{x}} - \hat{\mathbf{x}}_{\hat{S}}\|_F^2}{\|\mathbf{x}\|_F^2}$, where $\hat{\mathbf{x}}_{\hat{S}}$ indicates an image whose entries are same as $\hat{\mathbf{x}}$ over $\hat{S}$ and zero elsewhere. RE measures the sum of the squared amplitudes of the spuriously recovered impulses relative to the energy of the ground truth. Lower values of the RE metric indicate that the spurious impulses have smaller amplitudes, thereby guaranteeing reliable reconstruction.

The values of TPR, SNR_1, and RE obtained using the BSR algorithm are reported in Table 1, for different values of p and σ_h. We observe that the BSR algorithm accurately estimates the locations of the point sources, resulting in TPR $= 1.0$. The values of SNR_1 and RE are acceptable for most applications, even for a relatively larger spread and size of the Gaussian kernel. However, increasing the values of p and σ_h intensifies the blurring effect, thereby resulting in smaller values of SNR_1 and larger values of RE, as one would typically expect. Despite the deterioration in the performance in terms

Kernel parameters	TPR	SNR_1 (dB)	RE (dB)
$p = 5, \sigma_h = 2$	1.0	28.40	-26.57
$p = 7, \sigma_h = 3$	1.0	25.09	-17.92
$p = 9, \sigma_h = 4$	1.0	22.26	-15.64
$p = 11, \sigma_h = 5$	1.0	18.99	-15.13
$p = 13, \sigma_h = 6$	1.0	16.47	-16.86
$p = 15, \sigma_h = 7$	1.0	13.04	-14.86

Table 1. Performance evaluation of the BSR algorithm in estimating images containing point sources; the Gaussian kernel is of size $p \times p$ pixels and has a standard deviation of σ_h.

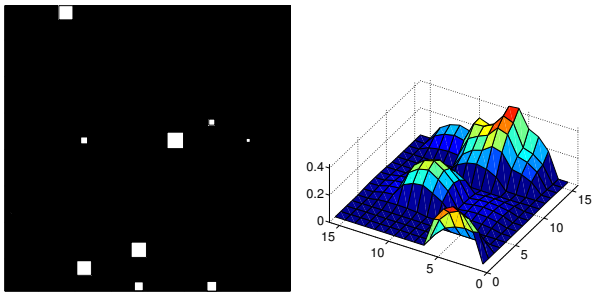

(a) Ground truth (b) Blurred measurement

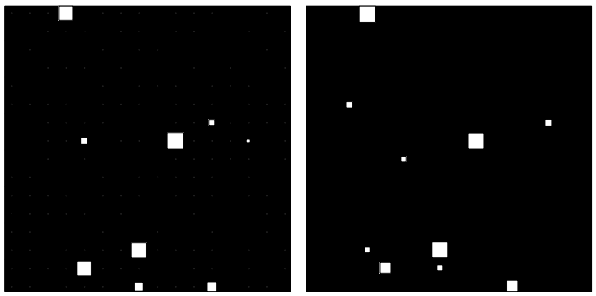

(c) Oracle recovery (d) BSR recovery

Fig. 2. (Color Online) Performance of BSR on a point-source image, consisting of $s = 10$ impulses, shown using Hinton diagrams. The locations and the amplitudes of the impulses are accurately recovered. The kernel parameters are chosen as $p = 5$ and $\sigma_h = 2$.

of SNR_1 and RE, with increase in p and σ_h, we observe that the BSR algorithm provides a reasonably good estimate of the locations of the impulses in the ground truth.

For visual comparison, we represent the ground truth and the images reconstructed using the BSR algorithm for a Gaussian blur kernel with $p = 5$ and $\sigma_h = 2$, in Fig. 2, using a Hinton diagram [25] that is quite popular in the *deep learning* community. The sizes of the squares in the diagram are proportional to the amplitudes of the respective impulses. The ground truth in the experiment consists of 10 point sources. We compare BSR with a deconvolution setup that acquires measurements without quantization, and the ground truth is recovered by solving

$$\hat{\mathbf{x}}_{\text{oracle}} = \arg\min_{\mathbf{x}} \ \|\mathbf{x}\|_1 \ \text{ subject to } \mathbf{y}_{\text{fp}} = AH\mathbf{x}, \qquad (8)$$

where $\mathbf{y}_{\text{fp}}$ denotes the full-precision measurement. The recovered signal using (8) is shown in Fig. 2(c) and indicated as the *oracle* recovery. We observe in Fig. 2(d) that the BSR algorithm recovers the

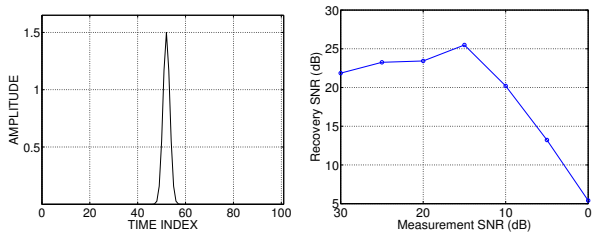

(a) Gaussian blur kernel (b) Recovery SNR versus input SNR

Fig. 3. Performance of the BSR algorithm in the presence of noise.

impulses with almost the same accuracy as the oracle, despite being supplied with binary measurements. This result is encouraging, as one can potentially leverage it to build fast and inexpensive systems for resolving point sources in a real-world imaging application.

3.3. Performance Evaluation in Noise

To conduct an empirical study of the BSR algorithm under measurement noise, we carry out experiments considering noisy measurements of the form $y_i = \text{sgn}\left((AH\mathbf{x} - \tau)_i + w_i\right)$, where w_i denotes the measurement noise. If $\mathcal{F}$ denotes the set measurements whose signs get flipped due to noise, we observe that the condition (1) is violated for every $i \in \mathcal{F}$, thereby rendering (5) infeasible. To circumvent this problem, we solve a relaxed version of (5), given by

$$\min_{\mathbf{x},\tau,\xi_i} \ \left\|\Lambda^{(p)}\mathbf{x}\right\|_1 + \beta \sum_i^m \xi_i \ \text{s.t.} \ y_i \left(\phi_i^T \mathbf{x} - \tau\right) \geqslant 1 - \xi_i, \ \xi_i \geqslant 0, \ (9)$$

where $\beta > 0$ is an appropriately chosen constant. The measurement consistency requirement is relaxed by introducing a set of non-negative slack variables ξ_i.

The experiment is performed on the superposition of point sources in 1-D, as shown in Fig. 1(a). The Gaussian blur kernel shown in Fig. 3(a) is used as the impulse response of the measurement device. The number of binary measurements acquired is $m = 2l_z = 600$. The number of iterations is chosen as $I = 8$. The reconstruction SNR, as a function of input/measurement SNR, is plotted in Fig. 3(b). We observe that the BSR algorithm is capable of tolerating noisy measurements up to an SNR of 15 dB. If the input SNR reduces below 15 dB, the drop in the reconstruction SNR is gradual and acceptable. We chose $\beta = 0.02$ in the experiment and found that it works reasonably well for a range of input SNR values. However, the choice of β for optimum recovery performance requires further investigation and will be addressed separately. As a rule of thumb, one should use larger (smaller) values of β for higher (lower) measurement SNR.

4. CONCLUSIONS

We propose an iterative convex programming-based algorithm, referred to as BSR, for recovering point sources from binary measurements of the random projections of its blurred counterpart. Our algorithm does not assume the knowledge of the number of impulses to be recovered and the threshold used for binary encoding. We have demonstrated, using 1-D and 2-D examples, that the BSR algorithm recovers the locations and the amplitudes of the point sources accurately, even under considerable blurring. The binary measurement scheme can potentially result in faster acquisition and lesser storage

requirement, without compromising on the performance. We have also developed a modification of the BSR algorithm to work with noisy measurements and demonstrated that it performs reliably up to a measurement SNR of 15 dB.

5. REFERENCES

[1] H. Greenspan, "Super-resolution in medical imaging," *Computer Journal*, vol. 52, issue 1, pp. 43–63, Aug. 2006.

[2] K. G. Puschmann and F. Kneer, "On super-resolution in astronomical imaging," *Astronomy and Astro-physics*, pp. 436–373, no. 378, 2005.

[3] T. A. Klar and S. W. Hell, "Subdiffraction resolution in far-field fluorescence microscopy," *Optics Letters* vol. 24, no. 14, pp. 954–956, 1999.

[4] B. Huang, W. Wang, M. Bates, and X. Zhuang, "Three-dimensional super-resolution imaging by stochastic optical reconstruction microscopy," *Science*, vol. 319, no. 5864, pp. 810–813, Feb. 2008.

[5] C. W. McCutchen, "Superresolution in microscopy and the Abbe resolution limit," *J. Opt. Soc. Am.*, vol. 57, no. 10, pp. 1190-1192, 1967.

[6] J. L. Miller, "Superresolution microscopy reveals chromosomes' smallest structure," *Physics Today*, vol. 68, no. 5, pp. 14–16, May 2015.

[7] E. J. Candès and C. Fernandez-Granda, "Towards a mathematical theory of super-resolution," *arXiv:1203.5871v3*, Nov. 2012.

[8] E. J. Candès and C. Fernandez-Granda, "Super-resolution from noisy data," *arXiv:1211.0290v3*, Jul. 2013.

[9] A Bourquard, F. Aguet, and M. Unser, "Optical Imaging Using Binary Sensors," *Optics Express*, vol. 18, no. 5, Mar. 2010.

[10] A Bourquard and M. Unser, "Binary compressed imaging," *IEEE Trans. Image Process.*, vol. 22, no. 3, pp. 1042–1055, Mar. 2013.

[11] P. T. Boufounos and R. G. Baraniuk, "1-bit Compressive Sensing," in *Proc. Conf. on Info. Science and Systems* (CISS), Princeton, NJ, Mar. 2008.

[12] A. Gupta, R. Nowak, and B. Recht, "Sample complexity for 1-bit compressed sensing and sparse classification," in *Proc. Int. Symposium on Info. Theory* (ISIT), 2010.

[13] Y. Plan and R. Vershynin, "One-bit compressed sensing by linear programming," *arXiv:1109.4299v5*, Mar. 2012.

[14] Y. Plan and R. Vershynin, "Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach," *IEEE Trans. Info. Theory*, vol. 59, no. 1, pp. 482-494, Jan. 2013.

[15] P. T. Boufounos, "Greedy sparse signal reconstruction from sign measurements," in *Proc. Asilomar Conf. on Signals, Systems, and Computation* (SSC), Asilomar, CA, Nov. 2009.

[16] M. Yan, Y. Yang, and S. Osher, "Robust 1-bit compressive sensing using adaptive outlier pursuit," *IEEE Trans. Signal Process.*, vol. 60, no. 7, pp. 3868–3875, Jul. 2012.

[17] L. Jacques, J. N. Laska, P. T. Boufounos, and R. G. Baraniuk, "Robust 1-Bit compressive sensing via binary stable embeddings of sparse vectors," *arXiv:1104.3160v4*, Nov. 2015.

[18] A. Mohaved, A. Panahi, and G. Durisi, "A robust RFPI-based 1-bit compressive sensing reconstruction algorithm," in *Proc. IEEE Info. Theory Workshop* (ITW), Sep. 2012.

[19] J. N. Laska and R. G. Baraniuk, "Regime change: Bit-depth versus measurement-rate in compressive sensing," *arXiv:1110.3450v1*, Oct. 2011.

[20] P. T. Boufounos, "Reconstruction of sparse signals from distorted randomized measurements," in *Proc. Intl. Conf. on Acoustics, Speech, and Signal Process.* (ICASSP), Mar. 2010.

[21] U. S. Kamilov, A. Bourquard, A. Amini, and M. Unser, "One-bit measurements with adaptive thresholds," *IEEE Signal Process. Lett.*, vol. 19, no. 10, pp. 607–610, Oct. 2012.

[22] J. Fang, Y. Shen, and H. Li, "One-bit quantization design and adaptive methods for compressed sensing," *arXiv:1304.1969v2*, May 2013.

[23] E. J. Candès, M. B. Wakin, and S. P. Boyd, "Enhancing sparsity by reweighted ℓ_1 minimization," *J. Fourier Anal. Appl.*, vol. 14, pp. 877–905, 2008.

[24] M. Grant and S. P. Boyd, "CVX: Matlab software for disciplined convex programming," version 2.0 beta, Sep. 2013.

[25] G. E. Hinton and T. Shallice, "Lesioning an attractor network: Investigations of acquired dyslexia," *Psychological Review*, vol. 98, pp. 74–92, 1991.